Revealed Wisdom of the Crowd: Bids Predict Loan Quality*

Jiayu Yao

Nanyang Business School, Nanyang Technological University

jiayu.yao@ntu.edu.sg

Mingfeng Lin

Scheller College of Business, Georgia Institute of Technology

mingfeng.lin@scheller.gatech.edu

D.J. Wu

Scheller College of Business, Georgia Institute of Technology

dj.wu@scheller.gatech.edu

October 2024

* All authors contributed equally. The authors thank Prabuddha De, Noyan Ilk, Karthik Kannan, Weifeng Li, Mohammad Rahman, and Xuan Wei for their helpful insights. We are grateful to conference participants at the Conference on Information Systems and Technology 2020, the INFORMS Annual Meeting 2020, the MISQ Author Development Workshop 2021, the International Conference on Smart Finance 2022, the Symposium on Statistical Challenges in Electronic Commerce Research 2023, and seminar participants at Chinese University of Hong Kong, Emory University, Georgia Institute of Technology, Lehigh University, Nanyang Technological University, Purdue University, University of British Columbia, University of Connecticut, University of Delaware, University of Georgia, University of Hong Kong, University of Illinois Urbana-Champaign, University of New Hampshire, University of Notre Dame, University of Texas at Austin, University of Texas at San Antonio for their comments and suggestions. We also gratefully acknowledge the financial support of the Business Analytics Center (BAC) in the Scheller College of Business and the Ewing Marion Kauffman foundation.

Revealed Wisdom of the Crowd: Bids Predict Loan Quality

Abstract

Despite the popularity of the phrase “wisdom of the crowd,” not all crowds are wise because not everyone in them acts in an informed, rational manner. Identifying informative actions, therefore, can help to isolate the truly “wise” part of a crowd. Motivated by this idea, we evaluate the informational value of investors’ bids using data from online, debt-based crowdfunding, where we were able to track both investment decisions and ultimate repayment statuses for individual loans. We propose several easily scalable variables derived from the heterogeneity of investors’ bids in terms of size and timing. We first show that loans funded with larger bids relative to the typical bid amount in the market, or to the bidder’s historical baseline, particularly early in the bidding period, are less likely to default. More importantly, we perform theory-driven feature engineering and find that these variables improve the predictive performance of state-of-the-art models that have been proposed in this context. Even during the fundraising process, these variables improve predictions of both funding likelihood and loan quality. We discuss the implications of these variables, including loan pricing in secondary markets, crowd wisdom in different market mechanisms, and financial inclusion. Crowdfunding platforms can easily implement these variables to improve market efficiency without compromising investor privacy.

Keywords: wisdom of the crowd, platform, crowdfunding, fintech

1. Introduction

The crowd is a fascinating phenomenon. A group of many individuals, each with their own intelligence, flaws, incentives, biases, and preferences, can collectively predict changes in the stock market, who will win the next election, and create an encyclopedia that eclipsed the professionally edited Britannica. However, the crowd is not always wise. While extreme events such as the Dutch Tulipomania of 1634-7 are not frequent, a casual look at the stock market and other online markets, especially financial technology (fintech) markets, reveals that crowds are neither always prudent nor always prescient. While one may be satisfied that crowds are more often correct than incorrect, there is clearly a need to better understand and unpack the actions of individuals in the crowd, particularly if it helps to identify a part of the crowd that is wiser than the rest. This would produce a deeper understanding of the boundaries of crowd wisdom and reduce the likelihood of future incidents of crowd irrationality.

The crowd is neither a monolithic nor a stable entity. Naturally, it has members who are better at prediction than others, either through access to better/more information or better ability to interpret the same public information, as well as those who occasionally blindly follow the actions of others (i.e., social learning). The actions of the first group should be more informative than the latter. Consequently, if we can distinguish between them, we may be able to derive higher levels of crowd intelligence. The challenge clearly lies in the fact that the smart crowd members are not easy to identify. They almost never have conspicuous labels attached to them, and their decision processes or judgments are rarely known. Existing studies have generally examined crowd heterogeneity in terms of demographic variables such as gender, which is static and often sensitive; or experience, which can be slow to change as investors learn from the outcome of past decisions. Our goal is to identify crowd heterogeneity based on information that is less sensitive than demographics, more timely than experience, and therefore more informative. Inspired by revealed preference theory, we infer investors' judgment through their actions.

More specifically, we propose several new variables derived from investor actions and show that they not only correlate significantly with investment quality but also improve the performance of predictive models. Since these variables are derived from revealed preference theories, we call them revealed wisdom of the crowd (RWOC) variables. The RWOC variables capture the relative amount of each bid compared to other bids on the market as well as to other bids from the same investor. They also capture the bid timing during the funding process.

We choose online debt crowdfunding (also known as peer-to-peer lending), particularly Prosper.com, as the empirical context to test the predictive value of these new variables. On these platforms, one can invest in fractions of loans posted by individual borrowers who are verified by the platform but remain anonymous on the site for privacy reasons. This context has many advantages over other types of crowdfunding. Most importantly, these loans are fixed-income assets with a predetermined repayment date, allowing us to objectively assess the wisdom of investing in them.

To this end, we follow the call of Shmueli and Koppius (2011) and Hofman et al. (2021) to integrate descriptive and predictive modeling and implement a theory-based feature-engineering approach. Specifically, we first propose a series of relationships between the RWOC variables and loan quality, and then test them using data from Prosper.com. Notably, we find that while the number of investors (size of a crowd) is not statistically associated with loan performance, this pattern is due to the mixed effects of signals and noise: The total number of large bids (compared to the rest of the market or an investor's own historical baseline) is significantly associated with a lower likelihood of default, especially when the large bids are placed early in the funding process. In contrast, the number of minimum bids (from investors who participate only to diversify their portfolio and may not carefully evaluate each loan) does not contain much useful information.

We then draw on the existing literature—including several papers that have used predictive models on data from Prosper.com (Fu et al. 2021, Iyer et al. 2016)—to build predictive models with RWOC variables as new features for predicting loan quality. We show that, after adding RWOC variables, the models improve by 21.19% in the Area Under the Receiver Operating Characteristic Curve (ROC-AUC). We further demonstrate that these performance improvements do not arise from alternative explanations such as the presence of expert investors. These results show that RWOC variables contain the truly informative part of crowd actions, or the real "wisdom of the crowd."

One concern is that our analyses require the full bidding history, which will only be available after a loan has been funded. To address this, we show that RWOC variables also have predictive efficacy during the funding process. RWOC variables based on early bids help predict funding likelihood and speed of a loan request, as well as loan quality if funded.

Our approach also provides useful insights into the issue of financial inclusiveness, which is an essential aspect of crowdfunding's societal value proposition. We show that the predictive improvement is larger for borrowers who tend to have less access to capital: women and those with lower credit grades. This means that if a marketplace can implement RWOC variables and make them available to investors, it should help bridge gaps in capital access.

Our study has both theoretical and practical implications. Theoretically, it highlights the importance of opening the black box of the crowd to examine individual members before aggregating their actions. Instead of generalizing that a crowd is wise or irrational, separating signal from noise in a crowd is important. Practically, our study shows that, first, we can still strike a balance between leveraging the crowd's wisdom and preserving investor privacy. Platforms can easily replicate our approach to help investors make better decisions without compromising their privacy. In fact, using RWOC variables represents a much more efficient way for investors to process information contained in the full bidding history and thereby improve market efficiency. Second, our approach suggests that for assets that have gone through market matching, that matching process contains valuable information that has been generally ignored to date. For assets with resale opportunities (such as the secondary market for loans), RWOC variables can help better price these assets on the secondary market. We conclude by discussing the generalizability of our findings under different market mechanisms and types of crowdfunding, as well as their implications for future research on participation thresholds on these online platforms and retail investor protection.

2. Empirical Context

We use data from debt-based crowdfunding, also known as "peer-to-peer lending" or P2P lending, which allows individuals to obtain unsecured personal loans from others. As of 2023, P2P lending has been the largest crowdfunding type, with a global market valued at $143.64 billion.[1] Given the market size, even tiny improvements in loan-quality prediction can be highly useful for stakeholders.

Data from debt crowdfunding offer unique advantages over prediction markets and other types of crowdfunding. Prediction markets typically focus on one-off events, which are unique and almost impossible to repeat (e.g., whether a candidate would win a certain election in a given year). In contrast, debt crowdfunding data offer many more comparable observations, enabling us to evaluate the wisdom of the crowd. Relative to donation- and rewards-based crowdfunding, in which donors and backers are likely to be at least partially driven by altruistic rather than economic motives (Dai and Zhang 2019, Hong et al. 2018, Simpson et al. 2021, Song et al. 2022), investors in debt crowdfunding have clearer incentives to evaluate loans and invest only in those that will preserve their funds and generate returns. Compared to equity crowdfunding in which investors are typically allowed to change their minds and withdraw their investments within a certain period (e.g., the "cooling off" period in Crowdcube) (Agrawal et al. 2014, Donovan 2021), debt crowdfunding typically does not allow bid withdrawal. Each bid, therefore, should represent more serious consideration and commitment and is more likely to be economically informative. Furthermore, the investment decision quality in debt crowdfunding is more straightforward to evaluate because when loans mature, one can unequivocally gauge whether each investor's decision to bid is economically justified. In contrast, firms in equity crowdfunding have an infinite time horizon, as they can be profitable one year and nearly bankrupt the next.

We use data from Prosper.com, one of the largest online P2P lending platforms in the United States. The Prosper transaction process involves several steps.[2] First, borrowers create accounts on the

[1] Numbers from Peer-to-Peer (P2P) Lending Global Market Report 2023 (https://www.researchandmarkets.com/reports/5766977).
[2] For more details about crowdfunding transactions, see Wang and Overby (2020) and Wei and Lin (2017).

website after verifying personal financial information. Second, they post a loan request, called a "listing," with the requested amount, interest rate[3], loan purpose, fundraising duration, and credit information about the loan and the borrower. Meanwhile, lenders can browse all available listings on the website and choose to bid on any they find interesting. In keeping with the literature (Iyer et al. 2016, Lin and Viswanathan 2016, Zhang and Liu 2012), "bid" in this context refers to an offer of an investment from a lender. An individual lender can bid as little as $25 (the minimum required by Prosper.com) or as much as the entire requested amount. Once placed, bids cannot be withdrawn. A listing closes successfully when 100% of the requested amount is received before expiration; otherwise, bids are refunded to lenders. Successful listings are further reviewed by Prosper.com staff members. Once approved, a listing becomes an issued loan. The borrower then receives the funds minus platform service fees. The borrower makes monthly repayments for the principal and interest, which the platform distributes proportionally to each investor. If a monthly payment is delayed by two or more months or the borrower stops payment, the account is sent to a collection agency. If the loan is still not fully repaid after the collection agency's efforts, it is treated as a loss. Lenders are explicitly informed that their investments are not guaranteed. The risk inherent in these lending decisions, therefore, provides another reason why this market serves as an interesting and useful context to study crowd decisions under uncertainty.

We highlight another two important Prosper features. First, borrowers cannot selectively repay some lenders but not others. Technically, all loans are made by WebBank, a Utah-chartered industrial bank. Borrowers obtain loans from the bank, while lenders receive loan shares. Second, during our study period, lenders could observe the bid history for each listing during the funding process: who bid how much, and when. Through the platform's public data export or application programming interface (API), lenders also had access to other raw data, including lenders' investment history and monthly loan payoffs. However, not all lenders would be willing or able to digest and/or properly interpret these data.

[3] Since 2011, Prosper.com has implemented a fixed interest rate model in which the platform determines the interest rate, and this rate does not change after the listing is posted.

3. Literature Review and RWOC Variables

Our study draws on and contributes to multiple research disciplines, including economics, behavioral finance, and information systems. In this section, we first review the literature on the wisdom of the crowd. Then we review how the literature has been applied in the context of crowdfunding since crowdfunding platforms ultimately depend on the crowd's collective intelligence to identify worthy investment opportunities. Finally, we build on recent studies of crowd heterogeneity, moving from individual-based to action-based differences to better extract crowd wisdom.

3.1 Wisdom of the Crowd

The key idea behind this notion is that a large group of nonexperts can be collectively more intelligent than a few experts (Surowiecki, 2005). While each individual's judgment may be biased, aggregate judgments tend to be surprisingly accurate, with errors associated with individual judgments canceling out. In other words, the key feature of this literature is the aggregate behavior of a crowd, rather than the uniqueness of each member within it.

A prime example of the wisdom of the crowd is crowdsourcing, which harnesses large networks of potential contributors to solve challenging problems. Not only has it received much scholarly attention (e.g., Greenstein and Zhu 2018, Lee et al. 2018, Lukyanenko et al. 2019, Majchrzak and Malhotra 2016, Wang et al. 2017), many companies and institutions have used it to collect ideas for new products and services. This is part of the "open innovation" movement and has been widely recognized as an approach to fostering innovation in an economical and efficient manner (Bayus 2013, Boudreau et al. 2011). The winning solution in crowdsourcing typically comes from a particular team or a member of the crowd. In other words, it only requires one team to come up with the best solution, so there is no need for a crowd consensus. In crowdsourcing contests, there is little emphasis on the solutions that were not chosen; information in the forsaken solutions is typically ignored.

Prediction markets and prediction polls are also valuable applications of crowd wisdom. In the former, where participants are motivated by potential profits, members buy or sell token stocks based on the information they possess or receive over time. The aggregated result, as reflected in the price of that "stock," often provides accurate predictions of future events (Chen et al. 2014, Da and Huang 2020, Spann and Skiera 2003). In the latter, forecasters are asked to assign probabilistic judgments of the likelihood of future events, either independently or interactively (e.g., engaging in discussions before making predictions) (Atanasov et al. 2017). What makes both different from crowdsourcing is that the wisdom of the crowd in prediction markets and polls is based on the *aggregate* decisions of its members, with each influencing the final result.

Crowdfunding is an especially interesting and complex example of the wisdom of the crowd due to its financial nature. Investors in crowdfunding have a strong incentive to fund "right" projects and weed out inferior ones. Each project must convince a sufficient number of backers to become successful. Many members of the crowd often participate repeatedly in the market, and each person's investment could be dramatically different (i.e., different weights in the total investment) based on their personal interpretations of the investment opportunity.

3.2 Wisdom of the Crowd in Crowdfunding

Crowdfunding has become an essential method of raising capital for both individuals and businesses. It harnesses the broad reach of the Internet to allow many investors to contribute to one individual or business. Along with the rapid growth of crowdfunding, academic research on this phenomenon has blossomed. Our goal is not to be exhaustive in reviewing this literature (Agrawal et al. 2014 and Moritz and Block 2016 offer excellent reviews) but to focus on research that directly informs our study.

A majority of research in crowdfunding, consistent with other literature on the wisdom of the crowd, examines how investors as a group interpret information about borrowers or fundraisers. Studies across many types of crowdfunding, including rewards-based, equity, and debt, suggest a remarkable level of wisdom in the financial decisions made by the crowd of small investors. In rewards crowdfunding, Mollick and Nanda (2016) find that crowd choices of projects on Kickstarter.com are similar to those of experts. In equity crowdfunding, Kim and Viswanathan (2019) have shown that early investors with work experience in related areas tend to select better projects; interestingly, the crowd can take advantage of that information by following their lead. In debt crowdfunding, where loan quality can be more objectively determined ex-post, there is also evidence that crowds perform remarkably well. Iyer et al. (2016) report that Prosper.com investors utilize multiple sources of information about borrowers, including unstructured and non-standard information, to predict their default likelihoods. This evidence suggests that the crowd, *as a whole*, is remarkably sophisticated.

There is also evidence of failures of collective intelligence. For example, Hildebrand et al. (2017) document adverse incentives in debt crowdfunding, finding that group leader bids are wrongly perceived as signals of high loan quality. Using data from LendingClub.com, Vallee and Zeng (2009) show that sophisticated investors perform better than retail investors, though this difference is attenuated when the platform makes less borrower information available. Leveraging a unique period of mispricing on Prosper.com, Lin et al. (2023) report that retail investors perform slightly worse than expert investors. Fu et al. (2021) propose a machine-learning algorithm that outperforms crowd predictions using data from Prosper.com.

Multiple studies have documented how investors react, correctly or incorrectly, to borrowers' demographic information, including race (Younkin and Kuppuswamy 2018), gender (Bapna and Ganco 2021, Ewens and Townsend 2020), appearance (Duarte et al. 2012), location (Agrawal et al. 2015, Burtch et al. 2014, Lin and Viswanathan 2016), as well as borrowers' social features such as social networks (Lin et al. 2013, Liu et al. 2015), social capital (Hasan et al. 2022), and text descriptions (Gao et al. 2023, Herzenstein et al. 2011). A common theme in these studies is that there exists a typical investor in the crowd who represents all investors, and the crowd is treated as a collective that interprets available information about borrowers. However, such quantifiable borrower characteristics cannot fully predict the quality of loans, a challenge also faced by commercial banks when issuing personal loans, despite the abundance of credit reports and other soft information. Precisely because of this challenge, not all investors would react to the same borrower information in the same way. For example, suppose investor A almost always invests the bare minimum required by the platform ($25) in each loan but invests $100 in borrower X. In that case, this increased investment indicates strong confidence in this particular borrower. In contrast, investor B might always invest $200 in each loan but only invest $100 in borrower Y. Even though both X and Y receive $100, do these investments indicate different opinions on behalf of the crowd, and can such a difference be captured systematically to make the market more efficient? We hypothesize that such differences can and should be captured to better predict loan quality and reduce information asymmetry. Our theoretical underpinning of the study is consistent with those of prediction markets: it is not only what is available and accessible to everyone, but more importantly, it is about how each person reacts – most likely in different ways – to the same observation information. Prediction markets exploit such differences in the interpretation of facts, and, in our context, each investor's bidding action is his/her interpretation of the risk/return tradeoff of lending to a borrower. To our knowledge, the literature published to date has not addressed this possibility.

3.3 Value of Action Heterogeneity

The idea that actions reflect preferences and opinions was first formalized in revealed preference theory (Samuelson 1938). For decades, marketers have carefully studied consumer purchase behaviors, even highly aggregated ones, to infer preferences and opinions (e.g., Berry et al. 1995). In the Internet age, practitioners and marketers pay close attention to online consumer behavior and then use this information to do market segmentation (referred to as "behavioral targeting"). This practice is now an industry standard (e.g., Chen and Stallaert 2014, Choi et al. 2020, Trusov et al. 2016).

Our study also draws on other streams of literature on heterogeneity. Specifically, it is highly consistent with long-standing findings in the open-innovation literature (i.e., crowdsourcing). Jeppesen and Lakhani (2010) championed the value of the broadcast search—that is, by sending a problem to a more diverse group of experts, an organization is more likely to find a solution for it. Due to the use of bidding in our context, our study is also informed by research on auctions, particularly those on strategic bidding (e.g., Bapna et al. 2003). While these auctions we study are different from those for selling merchandise, "jump bidding" or its converse, "sniping" behavior, documented in the literature suggests that large-size bids placed earlier (or later in the case of "sniping") reveal useful information that strategic bidders can exploit. This inspires us to consider both bid size and timing. Meanwhile, we contribute to the auction literature by using data from a context that allows us to directly link bidding behaviors to the ex-post quality of the product (loan repayment) to measure bid informational values.

The heterogeneity we focus on builds on, but also differs from, expertise heterogeneities recently studied. In those studies, experts within a crowd are nearly always better at evaluating identical objective information about fundraisers. For example, Kim and Viswanathan (2019) report on investor differences regarding their prior work experience in software development. Such differences are conspicuous and observable by other crowd members, and the experts are better at identifying higher-quality investments. The differences we examined are broader and not attached to specific individuals; we hypothesize that there can be economically meaningful information extracted from the actions of potentially *every* investor.

4. Descriptive Model of RWOC Variables and Loan Quality

The primary goal of our study is to identify new variables that improve predictive performance, and we implement a theory-driven feature-engineering approach.[4] Before constructing our models, we first follow the suggestions from Hofman et al. (2021) and Shmueli and Koppius (2011) to conduct descriptive analyses and provide initial evidence on the relevance of the RWOC variables. Specifically, we draw on the literature that leads us to propose the RWOC variables and investigate the relationships between them and loan quality using data from Prosper.com.

4.1 RWOC Variables

Investment actions can be characterized in two complementary ways: (A) amount (how much investors place their bids) and (B) timing (when investors place their bids) (Burtch et al. 2013, Grenadier and Wang 2005, Jiang et al. 2020). When an investor chooses an opportunity, their evaluation of it should be reflected in both dimensions. We therefore construct the RWOC variables from these dimensions.

[4] We thank an anonymous reviewer for proposing this term.

4.1.1 Investment Size

4.1.1.1 Absolute Investment Amount

On Prosper.com, investors can choose to invest partially in a loan for a minimum of $25, or to purchase an entire loan, which is obviously riskier due to lack of diversification. Consistent with the marketing literature on purchase quantity models (e.g., Krishnamurthi and Raj 1988), investment size reflects investors' assessments of the borrower's creditworthiness. A small investment suggests that the lender either does not find the borrower trustworthy or does not scrutinize the loan carefully—that is, the investor does not care about the specific loan enough to examine it closely, but simply trusts the platform's decision to list it, and is investing to diversify their portfolio (Herzenstein et al. 2011). Hence, minimum bids reveal very little information about an investor's evaluation of the loan beyond the static borrower information provided. In contrast, large investments, especially those significantly greater than the platform minimum, are more likely to reflect positive evaluations of the borrower's creditworthiness. Otherwise, the investor would have no incentive to place more funds at risk with a single loan. The relationship between bid amount and knowledge is well documented in the literature. Kim and Viswanathan (2019) show that expert investors with more knowledge place larger bids. All else equal, a higher-quality listing should be able to convince more investors to place larger bids[5]. Therefore, there should be a positive correlation between the number of large bids and the quality of the loan.

A potential counterargument against the above observation may arise from the nature of debt crowdfunding, where the loan amount (the debt that the borrower would take on) is determined prior to the start of the bidding process. This is different than equity- or reward-based crowdfunding where "overfunding" is possible and often embraced. Because of this, mathematically, a loan that attracts more

[5] As a quantitative example, assume that all investors have the same portfolio and budget and a minimum bid of $25, we compare two hypothetical loans that are both funded with $5,000 in total. The first loan (A) is funded by 200 investors with minimum bids of $25. The second loan (B) is funded by 5 investors at $1,000 each. This suggests that investors in Loan B demonstrate a much higher degree of confidence in their evaluations than investors in Loan A; therefore, Loan B is more likely to be a good investment because its investors would not have invested so much if they lacked confidence in its quality. Investors in Loan A, however, all simply meet the minimum investment threshold (which is more likely to be caused by a desire to diversify their portfolios), so their investment decisions neither reflect nor instill much confidence.

large bids would be funded by a *smaller number* of investors. This appears to contradict a common finding in the crowd wisdom literature (Mannes 2009) which says that a larger crowd is more likely to gather around something of a higher value. Under this logic, all else equal, there could also be a positive relationship between the number of bids (and equivalently, a smaller number of large bids) a listing receives and loan quality. Hence, the contrast between the number of larger bids and the size of the crowd attracted (larger number of smaller bids) represents an interesting tension driven by the fixed-loan-size feature of debt-crowdfunding that has not been previously addressed in the crowd wisdom literature.

4.1.1.2 Relative Investment Amount

We focus on absolute bid dollar values in the above. However, due to budget constraints and risk preferences, what one investor views as a large investment may be tiny for another. This heterogeneity means that investments of the same absolute dollar amount should reflect different confidence levels from different investors. If we track each investor's pattern over time, we can establish baselines to measure relative values.[6] Even when the bids contribute the same dollar amount to the loan, they can reflect dramatically different opinions. We thereby conclude that a bid contains significant informational value when it exceeds *a particular investor's typical investment amount* (henceforth, an "above-normal bid"). Regardless of the absolute amount, such bids indicate a positive evaluation that should be economically significant (Bower 2015). Evaluating bid informational value based on investors' historical behavior is logically consistent with Bansal and Gutierrez (2020) and Budescu and Chen (2015), who consider the historical performance of individual subjects before aggregating their forecasts. We therefore hypothesize that all else equal, there is a positive relationship between the number of above-normal bids a listing receives and loan quality.

[6] For example, for a wealthy investor with a million-dollar portfolio who typically invests $500 in each loan, a $500 bid may reflect little confidence and therefore does not carry too much useful information. In contrast, a $500 bid from a less-wealthy investor who almost always bids the platform minimum of $25 for other loans will reflect remarkably high confidence.

4.1.2 Investment Timing

Besides investment amounts, the other important aspect of investment decisions is their timing, which also reflects an investor's evaluation of loan quality. The norm in debt crowdfunding is that borrower information is available to potential investors at the beginning and does not change during the fundraising process.[7] If such information is sufficient to convince investors, they should be motivated to invest immediately, because loan values are fixed: Once a listing reaches its required funding amount, no further investments are possible. Investment opportunities in loans are limited, which is different from rewards crowdfunding where overfunding is possible. Thus, there is no upside for lenders who are confident in the listing to wait, as they might miss the opportunity. This is consistent with the psychology and decision science literature; for example, subjects' response time to a stimulus reveals the strength of preference or belief beyond that revealed by choice outcomes (Frydman and Krajbich 2022, Konovalov and Krajbich 2017). The finance literature also documents that investors with sufficient confidence or experience in evaluating borrower quality are more likely to bid early and rapidly (e.g., Jiang et al. 2020, Kim and Viswanathan 2019, Vallee and Zeng 2019). The model proposed by Holden et al. (2022) suggests that early lenders are more likely to be informed lenders. Following the logic of absolute investment amount, we expect to observe more large bids in the *earlier* stage of the funding process for higher-quality loans.[8]

Despite the argument presented above, there may still be incentives for investors to delay placing bids, even if they feel confident about the loan. Late investors can receive more cues from peer investors (e.g., social learning; Bikhchandani et al. 2021). Prior papers also document a rational herding effect in P2P lending (Zhang and Liu 2012). Lenders may benefit from confirming their independent evaluations by observing peer behavior and therefore strategically delay their own bids (Burtch et al. 2016, Jiang et al. 2022, Kim and Viswanathan 2019, Liu et al. 2015, Müller-Trede et al. 2018). In this scenario, it seems

[7] In rewards and equity crowdfunding, fundraisers can often respond to questions from potential funders or make changes during the fundraising process; this is generally not the case in debt crowdfunding.

[8] To use a more quantitative example, suppose there are two loans (C and D) that are otherwise similar—including the individual bids—with the only difference in bid sequence: C and D have the same loan amount and both were funded by one large ($1000) and 20 small ($25) bids. The only difference is that the $1000 bid began the bidding for C, whereas for D the $1000 completed the bidding. H3a predicts that C will be of higher quality than D.

plausible that more large bids are placed *later* in the funding process for higher-quality loans. However, these two observations do not directly contradict each other: The first one is that more large bids earlier in the bidding process is more likely associated with a good loan than loans with fewer large bids earlier in the process; the latter observation is that more large bids later in the bidding process is more likely associated with a good loan, than loans with fewer large bids later in the bidding process. Nonetheless, it is more difficult to attribute larger bids placed later in the bidding process to an investor's own insight because it could have been driven by the actions of peers. Earlier large bids are more likely to reveal independent, confident evaluations of investors, and therefore more likely to be indicative of loan quality.

4.2 Empirical Strategy

To test the relationship between bid heterogeneity that we propose above and loan quality, we sequentially classify bids as shown in Figure 1. The top level (Level 0) represents the total number of bids in a loan. Level 1 separates bids that are equal to Prosper's minimum ($25) from those above the minimum. We call the latter "above-minimum bids." Level 2 further classifies the above-minimum bids as moderate or large relative to all other bids in the market around the same time (details in §4.3). On Level 3, we further distinguish bids using individual lenders' bidding histories, i.e., whether the moderate/large bids were typical or atypical. Finally, on Level 4, we focus on bid timing to subdivide each Level 3 class into early and late. Importantly, on each level, the total numbers of all bid types sum up to the total number of bids for the loan; each level indicates different ways of classifying these bids. To test the relationships specified in our hypotheses, we sequentially test the relationships between bid-type variables and loan quality. Details follow our variable descriptions.

[Insert Figure 1 about here.]

4.3 Data and Variables

Our raw dataset includes all loan requests (funded and unfunded) posted on Prosper.com between January 2011 and December 2012, information available to lenders about borrowers and loan requests, bids placed on those loans (including their timing and amount as well as the investor who placed it), and the ex-post monthly repayment histories of funded loans. We selected this time window because the primary features on Prosper.com were highly consistent and stable throughout the period.[9] In our primary analysis, we focus on funded loans since unsuccessful listings do not become loans and therefore do not have objective quality data (see §2).

Our dataset excludes loans wholly funded by a single lender, as such loans have no lender heterogeneity data; there are only 573 such loans in the dataset (less than 2%). Below we discuss the primary variables of interest, each defined at the loan level of analysis. Online Appendix I provides definitions of all variables.

Loan quality: A key benefit of studying debt crowdfunding is that loans have an unequivocal measure of quality at their maturity. Since they are repaid monthly, we calculate the outcome as the proportion of principal that the borrower fails to repay (i.e., the ratio of the default amount to the total loan amount). We call this metric "loan default percentage" for consistency with the literature (e.g., Iyer et al. 2016, Lin and Viswanathan 2016). This value is zero if a loan is fully repaid, and lower percentages suggest higher quality. The robustness checks in §4.5 provide additional results from alternative outcome variables (i.e., a binary default indicator, and return on investment).

Absolute bid amount: Levels 1 and 2 differentiate bids by dollar amount. If the bid is equal to the minimum required by Prosper.com ($25), we consider it a *minimum bid*. We count the number of minimum bids for each loan to generate the variable *#minimum bids.* The difference between total bids and *#minimum bids* is defined as *#Above-minimum bids.* For Level 2, to differentiate between moderate and large bids, we calculate the mean and the standard deviation for all bids placed in the entire market during the week of the bid. If the bid amount is larger than the sum of the mean and 3 times the standard deviation, we define it as a *large bid* (also known as the "three-sigma rule"). Over 2% of lenders made such unusual bids at least once in our sample period. Otherwise, if a bid is larger than the minimum but

[9] Prosper.com shifted from an auction model to a fixed-rate model on December 19th, 2010. Beginning in January 2013, Prosper.com experienced a series of management team changes, new investments, and feature modifications. We therefore limit the time window to the 24 months between January 2011 and December 2012. During our study period, the platform used posted-price format only.

smaller than the mean plus 3 times the standard deviations[10], we define it as a *moderate bid*. We count the total number of moderate and large bids for each loan to generate *#Moderate bids* and *#Large bids*.

Relative amount: On Level 3, we compare each bid with the mode of all bids from the same lender. We treat each investor's *mode* as the typical bid that they place on loans.[11] If a bid is less than or equal to the mode, we define it as a *normal bid*; otherwise, it is an *above-normal bid*. We count such bids for each loan to generate *#Above-normal bids*.

Investment timing: We split the funding process into *early* and *late* halves of loan funding duration. Level 4 of Figure 1 shows how all bid types on Level 3 are temporally divided.

Control variables: We incorporate all available loan and borrower characteristics provided by Prosper for lenders, such as requested amount, interest rate, term, category, campaign duration, credit grade, income level, employment status, and home ownership. In addition, we have generated some derived features, such as description length and quarterly fixed effects.

Table 1 presents summary statistics. Our sample contains 29,900 loans funded by 2,494,942 bids. The average loan size is $7,458 (range $2,000 to $25,000). All loans in the dataset include at least one above-minimum bid, with 15,375 (51.42%) including one or more large bids. Those with at least one large bid receive $9,060 and 69 bids on average, while those without any large bids receive $5,863 and 98 bids on average. In terms of the relative bids, 60.92% of the large bids are larger than the lenders' mode, while 42.11% of the moderate bids are larger than the mode. In other words, not all large bids are unusual for the lenders who placed them, highlighting the value of relative bid amounts. Finally, 33.24 % and 66.76% of the bids are placed in the early and late stages, respectively.

[Insert Table 1 about here.]

[10] Results are highly consistent when we use the mean plus two standard deviations.
[11] Results are highly consistent when we use the median investment amount.

4.4 Models and Results

Before we delve into predictive models, we first present the methods, primary results, and robustness checks for our descriptive model that tests the relationship between RWOC variables and loan quality. To examine how different bid types predict loan quality, we estimate the following specifications:

$$DefaultPercentage_i = \beta \times BidTypes_i^L + \text{f}(ListingInfo_i + BorrowerInfo_i + QuarterlyFEs_i) \quad (1)$$

where $BidTypes_i^L$ represents all the bid types on Level L of Figure 1 for Loan i; $L \in \{0,1,2,3,4\}$, corresponding to the five levels in Figure 1.[12] β is a vector of coefficients for different bid types, as shown in Figure 1; when negative, a particular type of bid is associated with a lower default percentage and therefore higher loan quality. Since our outcome variable of interest in the primary specification is *Default Percentage*, which ranges between 0 and 1, we use a fractional outcome regression model (Lin and Viswanathan 2016, Pang 2017, Papke and Wooldridge 1996) to estimate the parameters of interest.

Table 2 reports our estimates. Each column reports the results of one of the bid hierarchy levels shown in Figure 1. Column 1 shows a non-significant relationship between the total number of bids in a loan and its default percentage. This suggests that when we do not differentiate among bids, the number of bids (in a sense, the size of the crowd) that a loan receives does not provide extra information about loan quality beyond what has already been conveyed in the borrower and loan variables. In other words, bid numbers, a traditional measure of popularity, fail to offer informational value in the context of debt crowdfunding, where overfunding is not possible.

More interesting patterns emerge as we move down the hierarchy. Column 2 (Level 1) shows that *#Minimum Bids* is positively correlated with the default percentage, whereas *#Above-minimum bids* is negatively correlated with the default percentage. In other words, of the bids that a loan receives, only those larger than the platform-mandated minimum are predictive of high loan quality.

---

[12] For example, all the bid types on Level 3 are Minimum bids, Normal moderate bids, Above-normal moderate bids, Normal large bids, and Above-normal large bids. Then the specification for examining Level 3 is $DefaultPercentage_i = \beta_1 \times \#Minimum\ Bids_i + \beta_2 \times \#Normal\ Moderate\ Bids_i + \beta_3 \times \#Above\ Normal\ ModerateBids_i + \beta_4 \times \#Normal\ Large\ Bids_i + \beta_5 \times \#Above\ Normal\ Large\ Bids_i + \text{f}(ListingInfo_i + BorrowerInfo_i + QuarterlyFEs_i)$.

Column 3 (Level 2) further categorizes above-minimum bids between moderate and large bids. The coefficient for *#Minimum Bids* remains positive and significant. *#Moderate Bids* and *#Large Bids* are both negatively associated with default percentage, consistent with the result in Level 1. However, the coefficient (and the related marginal effect) of *#Large Bids* is much larger than that of *#Moderate Bids*. Holding everything else at the median, each additional large bid produces a 0.61% reduction in default percentage, approximately $33 in default amount. These results lead us to formulate our first observation:

Observation 1: All else equal, there is a positive relationship between the number of large bids a listing receives and the resulting loan's quality.

Level 3 in Column 4 further distinguishes bids by their *relative* values, i.e., whether they are greater than the mode of that lender's bids. Within both moderate and large bids, only above-normal bids are significantly correlated with a lower default percentage. In contrast, none of the normal bids (including minimum bids) has a statistically significant coefficient. Holding everything else at the median, each additional above-normal large bid produces a 0.84% reduction in default percentage, approximately $45 in default amount. Additionally, normal and above-normal large bids both have larger negative coefficients (in absolute value) than the corresponding types of moderate bids, consistent with our findings in Level 2. These results lead to our second observation:

Observation 2: All else equal, there is a positive relationship between the number of above-normal bids a listing receives, and the resulting loan's quality.

With respect to investment timing, Level 4 categorizes all bid types in Level 3 into early and late stages. Most coefficient estimates for above-normal bids remain negative and significant (except the late above-normal large bids), but all such relations are stronger when the bids are placed in the early rather than the late half. For example, the coefficient for early above-normal large bids is more than two-fold greater than the coefficient for late above-normal large bids. Our calculation of the marginal effect shows that holding all else at the median, each additional early above-normal large bid results in a 1.25% reduction in default percentage, approximately $68 in default amount. These results lead to our third observation:

Observation 3: All else equal, there is a positive relationship between the number of early large bids a listing receives, and the resulting loan's quality.

[Insert Table 2 about here.]

4.5 Robustness of the Descriptive Model

While the focus of our study is the value of the RWOC variables for predictive purposes, we still conduct an extensive set of robustness tests for our descriptive models' findings to rule out alternative explanations. First, various potential sources for endogeneity are unlikely to be a first-order concern in our study. (A) Reverse causality (i.e., borrower repayment affects bid distribution) is highly unlikely in our context because borrower repayment and bid distribution formed during the funding period are temporally disjointed. When the pre-set, 14-day funding period ends, Prosper.com staff begin to verify the loan. The repayment period (36 months) begins only after the verification is complete. As prior studies of the same market indicate, predicting loan quality is a non-trivial task (Iyer et al. 2016), even for institutional investors (Lin et al. 2023). Hence, lenders cannot fully foresee repayment and then decide whether or how much to invest. (B) Spurious correlation (i.e., an unobserved error term that correlates with bid distribution as well as borrower repayment) would be plausible if borrowers could observe bid distributions during bidding and then decide whether and how much to repay. This concern is also highly unlikely in our context because the repayment process of loans is entirely managed by the platform (in collaboration with WebBank). As such, borrowers cannot selectively repay some lenders but not others; rather, repayment is proportionally issued to each lender according to their proportion of the loan.

We further conduct a series of robustness tests (see details in Online Appendix II). We use matching as an alternative method to verify our primary findings; rule out herding as an alternative explanation; rule out the possibility that borrowers repay due to lenders being friends and family or being local lenders; and rule out the possibility that the results are driven by expert (institutional) investors. Our results also hold when we use alternative measures of loan quality, including a binary indicator of loan default and the loan's overall return on investment; when we measure bid size using the percentage of bid amount; when we differentiate moderate and large bids based on two alternative thresholds; and when we redefine early bids as the bids placed on the first day or first two days.

5. Using RWOC Variables for Loan Quality Prediction

The descriptive analyses conducted above reflect the utility of RWOC variables and their within-sample predictive power. As our primary interest is whether incorporating these variables can help predict loan performance, this section demonstrates that including RWOC variables in the feature set improves predictive performance across a wide range of advanced prediction algorithms, as well as in published models that also seek to predict loan quality using Prosper.com data.

5.1 Predictive Analysis

Before delving into the details of the formal prediction models, we begin by investigating some model-free evidence to gauge the predictive power of RWOC variables. Specifically, we examine and compare the presence of RWOC variables in both paid and default loans. Our findings reveal that paid loans, in contrast to defaulted ones, tend to include more informative bids (e.g., large bids) and fewer “noise” bids (e.g., early minimum bids). In other words, RWOC variables show the potential to differentiate between paid and defaulted loans, reflecting the idea of classification modeling. These observations are consistent with and lend support to our findings in the descriptive models, highlighting the potential significance of RWOC variables in predicting loan quality. Next, we introduce the setup of the prediction models and discuss the results.

5.1.1 Model Setup

In this section, we train a series of standard, supervised, machine-learning models to predict loan quality. The outcome variable of interest is a binary *Default Indicator*, one if the loan defaults and zero otherwise (Fu et al. 2021, Iyer et al. 2016).

Our goal is to study whether we can improve the performance of loan-quality prediction by incorporating RWOC variables as additional predictors. We first define three baseline feature sets. The first feature set uses Consumer Credit Score and Prosper Score (henceforth Feature Set A) since they reflect the comprehensive evaluation of borrowers from consumer credit rating agencies and Prosper respectively. The second feature set includes all the observable borrower and loan characteristics on Prosper (henceforth Feature Set B). This baseline is consistent with the typical underwriting practice of basing loan decisions exclusively on borrower characteristics. Online Appendix I Panel C lists all the variables. Furthermore, we build a more demanding baseline that adds the total number of bids on a loan to Feature Set B, which corresponds to Level 0 in Figure 1 (henceforth Feature Set C). Feature Set C reflects the traditional belief of the wisdom of the crowd, taking the crowd as a whole and not considering differences among members.

We then construct a series of feature sets that incorporate RWOC variables. We first construct Feature Set D that adds the simple RWOC variables we use in the descriptive models (all bid variables on Level 4 of Figure 1 since Level 4 is the most comprehensive) to Feature Set B. This feature set allows us to evaluate the predictive power of intuitive RWOC variables.

To further explore the information contained in RWOC variables, we engage in feature engineering, a process that allows us to generate and transform more relevant variables from the raw data (Zheng and Casari 2018). Focusing on the three dimensions of RWOC (absolute bid amount, relative bid amount, and bid timing), we engineer approximately 300 features that represent bid heterogeneity along these dimensions, both mechanically and manually. This leads to Feature Set E, which incorporates absolute investment amount-related features; Feature Set F, adding relative investment amount-related features; and Feature Set G, with investment timing-related features. Finally, we construct Feature Set H that encompasses all the engineered RWOC features. By comparing the performance of models using baseline feature sets and those using RWOC feature sets under different algorithms, we can evaluate whether the RWOC variables improve loan quality prediction beyond what has been captured and fixed in borrower information known before the start of bidding.

We utilize four popular classification algorithms: the basic classification method, multinomial logistic regression (Logistic) (Cox 1958); two widely applied boosting methods, Light Gradient Boosting Machine (LightGBM) (Ke et al. 2017) and eXtreme Gradient Boosting (XGBoost) (Chen and Guestrin

2016); and a deep-learning method, multilayer perceptron (MLP) (Haykin 1994). We use the same dataset used in the descriptive models (loans funded between January 2011 and December 2012) to train and test the predictive models. To reduce evaluation bias, we repeat five-fold cross-validation twice, using different random seeds at each iteration (James et al. 2013). The ideal hyperparameters of models are identified by hyperparameter tuning with the grid search method. The entire predictive modeling process is performed in Python using scikit-learn (Pedregosa et al. 2011), xgboost (Chen and Guestrin 2016), lightgbm (Ke et al. 2017), and tensorflow (Abadi et al. 2016) libraries.

It is worth noting that machine-learning methods are highly variable with different algorithms and datasets. Hence, although our model does perform considerably better than those reported in existing literature (§5.2), our primary goal is to evaluate whether adding RWOC variables to models trained on the same dataset can improve the model's predictive performance (Fu et al. 2021, Song et al. 2022).

5.1.2 Results

Table 3 presents the predictive performance of multiple algorithms using various feature sets. Following the literature (Padmanabhan et al. 2022), we evaluate two performance metrics. The first is the area under the receiver operating characteristic curve (ROC-AUC), which is a widely applied measure of predictive performance. ROC-AUC is particularly useful to represent how well a model can distinguish between classes (default or not) when there is no obvious cutoff for a probability from which one can consider an observation belongs to a class (Fu et al. 2021, Iyer et al. 2016). The AUC ranges from 0 to 1, and a higher AUC corresponds with better predictive performance. The second metric is the financial impact of improving predictive power produced by adding RWOC variables (Padmanabhan et al. 2022). Since the models aim to correctly classify repaid and defaulted loans, we measure financial impact by calculating investment losses from misclassified loans (Wang et al. 2017). Misclassification can be either false positive (classifying a repaid loan as defaulted) or false negative (classifying defaulted as repaid). Since the misclassification costs of these two cases are asymmetric, we follow Simester et al. (2020) to calculate the loss from a false positive as the loss of its interest, and the loss from a false negative as the loss of its principal. Table 3 reports the results.

[Insert Table 3 about here.]

Comparing the three baseline models in each algorithm, we first find that the credit score alone (Feature Set A) can achieve as much as 58.08% predictive power of all financial features (Feature Set B)[13]. By comparing predictions with Feature Sets B and C, we find that if we only add total bids as an additional predictor, the performance does not significantly increase, and sometimes slightly decrease, AUC. This result is consistent with our descriptive analyses and supports our argument that the size of the investor crowd that a loan attracts is not useful in predicting the quality of the loan.

We next examine the predictive performance of models with RWOC variables and find that incorporating RWOC variables consistently improves predictive performance. The feature sets with RWOC variables (Feature Sets D – H) consistently outperform the corresponding baseline feature sets (Features Sets A-C), regardless of which algorithm is used (including the best performer in the baseline feature sets, MLP). Within every algorithm, Feature Set H with all engineered RWOC variables always performs the best. To calibrate the AUC improvement, we follow Iyer et al. (2016), who also published a series of models using Prosper.com data. Applying their procedure (Iyer et al. 2016, p1565, footnote 17) to our context, we find that adding only simple RWOC variables (Feature Set D) can increase AUC by as much as 8.86% ((0.7384 - 0.5) / (0.7190 - 0.5) - 1 based on Logistic) from Feature Set B. The inclusion of engineered RWOC variables improves predictive performance by as much as 21.19% ((0.7654 - 0.5) / (0.7190 - 0.5) - 1 based on Logistic). Even using the algorithm that performs the best in baseline models (XGBoost using Feature Set B), the inclusion of engineered RWOC variables improves predictive performance by 7.74% ((0.7840 - 0.5) / (0.7636 - 0.5) – 1). All results together lend strong support to the value of RWOC variables in predicting loan defaults. If we compare predictions against the more demanding baseline, Feature Set C, including RWOC variables still improves performance across all

[13] This is achieved by Logistic algorithm. We followed Iyer et al. (2016) and calculate it as (0.6272 - 0.5) / (0.7190 - 0.5) = 58.08%.

these algorithms. This result is particularly impressive given the tremendous effort and expenditure involved in collecting the financial information of borrowers.

Moreover, to test whether the performance of the feature sets that incorporate RWOC variables statistically improves from baseline features, we again follow the literature (Ben-Assuli and Padman 2020, Berg et al. 2020, Iyer et al. 2016) and examine the significance of the AUC increases. This approach was first presented in DeLong et al. (1988), and we implement it using the roccomp command in the STATA, as implemented in Iyer et al. (2016). We find that all algorithms' ROC-AUC improvements of Feature Set H from Feature Set B are significant at the 1% level, providing further support for the predictive power of RWOC variables.

Table 3 Column 2 shows that including RWOC variables results in a significant reduction in investment losses. For our sample of 29,900 loans, there are additional savings of as much as $1,044,571 (0.47% of the total funded amount) of Feature Set H. Even for the algorithm that generates the smallest misclassification costs for Feature Set B (XGBoost), the savings are still $638,002 (0.28% of the total funded amount). Given the magnitude of debt-based crowdfunding markets (over $25 billion funded globally as of 2020[14]), such improvement is substantial and economically meaningful.

To further gauge the informativeness of RWOC variables, we construct additional predictor sets and compare their performance. Considering scenarios where investor history is unavailable (e.g., due to data absence or privacy concerns), we construct Feature Set I, which excludes private lender information. It adds the RWOC variables of absolute bid amount and bid timing, but not relative bid amount. Compared with Feature Set B, Feature Set I, which is superior from a privacy point of view, also achieves a decent improvement (20.23% higher AUC).

We next follow the literature and construct Feature Set J with only engineered RWOC variables but no finance variables (Berg et al. 2020, Netzer et al. 2019, and Wang et al. 2021). Although RWOC variables alone will be rarely used in reality, we report its results for technical demonstration and

[14] https://blog.fundly.com/crowdfunding-statistics.

completeness, as it provides a baseline for comparison and demonstrates the effectiveness of our proposed method. A comparison of Feature Set J and Feature Set B shows that RWOC variables alone provide as much as 89.81% predictive power compared with what all financial information does. This confirms the importance of lender heterogeneous behaviors when borrower information is completely unavailable.

5.2 Comparisons with Existing Models

We next apply RWOC variables to previously published models in the literature to verify that the RWOC variables remain valuable in those models. Specifically, we replicate our baseline models using the datasets from two prior papers, Fu et al. (2021) and Iyer et al. (2016). Both studies employ the same data source (Prosper.com) and performance metric (ROC-AUC). We first reconstruct their datasets, which cover different time periods from ours. We then retrain the algorithms using Feature Set B on each dataset separately (Table 4). Across all algorithms, our baseline models generate higher AUCs than the published predictive models. Thus, the results demonstrate that our baseline models not only match but indeed outperform those from the aforementioned studies, further underscoring the value of incorporating the RWOC variables.

[Insert Table 4 about here.]

5.3 Performance Improvements Do Not Come from Experts

Similar to our robustness checks in the descriptive analysis, we examine whether predictive improvements provided by the RWOC variables are fully contributed by the bids from expert investors. Expert investors tend to perform better than normal investors (Jiang et al. 2020, Kim and Viswanathan 2019, Lin et al. 2023). If expert bids have better predictive power than others, the bidding behaviors of expert lenders should be separately considered in our models.

We first construct a baseline with the financial variables and all the bid types on Level 2 of Figure 1 (Feature Set K). Unlike our primary predictive models that use Level 4 bids as RWOC variables, here we use Level 2 bids. This is because, as we move down through the hierarchy, some types of bids in Level 3 or Level 4 do not contain any bids from experts. To ensure sufficient data to train the models, we use bid types on Level 2. Next, we train three models that replace RWOC variables with three bid sets:

Level 2 bids from only experts (Feature Set L), Level 2 bids from only non-experts (Feature Set M), and Level 2 bids from both experts and non-experts (Feature Set N). Table 5 presents ROC-AUC from different algorithms (columns) using different feature sets (rows).

[Insert Table 5 about here.]

Comparing bids from experts with those from non-experts, we find that the latter (Feature Set M) has higher predictive power than the former (Feature Set L). In other words, the increase of predictive power provided by adding RWOC variables is primarily contributed by retail investors rather than institutional investors. Interestingly, across all algorithms, Feature Set N does not improve further, and even performs slightly worse, than Feature Set K, indicating that further dividing *action* variables by *identities* does not generate extra value. This result again confirms the value of the crowd-based RWOC variables.

5.4 Platform Auction Format and Predictive Model Performance

The predictive model results reported above are based on data from Prosper after it implemented the fixed-price format, where the platform pre-determines the interest rate of a loan and lenders decide whether to participate under this interest rate. A natural question is whether our findings regarding RWOC apply under the auction model. We therefore replicate the predictive model using Prosper.com data before this regime change, i.e., when loans were funded through auctions (Wei and Lin 2017). We find the predictive model performs even better under auctions. Even after we match the loans from these two formats based on interest rates, borrower characteristics, and other loan characteristics, the predictive model still performs better under auctions. This is consistent with Wei and Lin (2017) in that under the fixed-price format, the platform tends to price the loans' interest rate higher than the crowd would do to attract investors, thereby distorting the market and leading to efficiency loss.

6. RWOC Variables from Early Bids

While the above analyses demonstrate the value of adding the RWOC variables, one may argue that these are only achievable upon observing the full sequence of lender investments after the funding process has ended. While this would be useful for secondary market pricing, it would be even better if RWOC variables could help investors avoid investing in a bad loan in the first place. We therefore test the value of RWOC variables *during* the funding process, by only using early bids to predict quality, funding likelihood, and funding speed.

6.1 Predicting Loan Quality Using Early Bids

We start our analyses by investigating whether early bids alone predict loan quality using the same tests described in §5.1. Similar to the previous models, we build three revised models, in which the predictors include early bids that are separately defined as bids placed in the first half of the total funding time, the same definition used in the descriptive analyses (Feature Set O); bids placed on the first funding day (Feature Set P); and bids placed on the first two funding days (Feature Set Q) separately. Table 6 Column 1 reports the results. Comparing Feature Set O with previous feature sets, the default was predicted with as much as 18.22% greater accuracy than Feature Set B and achieves as much as 97.99% of the increase achieved by the most comprehensive feature set, Feature Set H. Feature Sets P and Q also achieve similar improvements from the baseline models. We also supplement our analysis by implementing an alternative training and test method, mirroring the approach detailed in Benjamin and Raghu (2023). In this method, we train the predictive models using the full bidding history (Feature Set H) but test using only features generated from early bids (Feature Set O, P, and Q). This method leverages all available data during the training phase and allows fast real-time predictions in live environments. The outcomes produced by this method (not displayed in the interests of space) are qualitatively consistent with our main results. Therefore, we show that early bids alone, though not as powerful as the entire set of RWOC variables, improve predictive performance significantly. Considering the fact that early bids account for only 37.25% of total bids in our dataset, such an improvement further attests to the value of the RWOC approach.

[Insert Table 6 about here.]

To further demonstrate the predictive power of early bids, we simulate two real-world detection scenarios. Since bids are timestamped, we examine how predictive power changes as bidding progresses.

We first create a series of separate feature sets for every additional decile of cumulative funding (i.e., 0%, 10%, …, 90%, 100%; 11 in total). Each feature set is composed of all finance variables and all the bids that have been placed till the decile. Figure 2 Panel (A) plots AUC as a function of funding percentage. For brevity, we present only the LightGBM algorithm, though the other algorithms are qualitatively consistent. The figure shows that, during early bidding, the predictive power of the RWOC variables increases rapidly when more early bids are placed. However, the AUC increase slows down at a funding percentage of approximately 40%, after achieving over 70% of the total increase. In other words, RWOC variables significantly improve predictions even if only bids from the early funding period are used.

[Insert Figure 2 about here.]

Similarly, we then create another series of separate feature sets for different funding periods (i.e., 24h, 48h, 72h, etc.) and plot the relation between AUC and funding time in Figure 2 Panel (B). We find a similar result as above: AUC increases with the funding period but slows down later. Our findings in this section are consistent with what we find in the descriptive model—compared with late bids, early bids contain more useful information and are more strongly associated with loan quality.

6.2 Predicting Full Funding Using Early Bids

In addition to the prediction of loan quality, another important application of early bids, as prior studies have addressed (Burtch et al. 2016, Zhang and Liu 2012), is to predict the funding success of listings. Compared with borrower information available before funding starts, early RWOC bids partially reveal market sentiments, helping to predict funding success. Since we are now examining full funding likelihood, we expand our sample to include all listings that were not funded in our study period (between January 2011 and December 2012). Before building a predictive model, we perform a series of survival analyses using listing-day level panel data and verify that there is indeed a positive relationship between early large bids and completing funding more rapidly (see Online Appendix III for more details).

We examine the power of early bids in predicting funding success under the prediction framework. We build revised models using three definitions of early bids (Feature Sets O, P, and Q) and predict a binary outcome variable for funding success—the full funding indicator, which is one if a listing is fully funded and zero otherwise. Table 6 Column 2 displays the results for each algorithm using different feature sets. Across all the algorithms, feature sets with early bids (Feature Set O, P, and Q) consistently outperform the corresponding baselines without information from early bids (Feature Set B). After adding early bids to finance variables, AUC increases as much as 76.10% for Feature Set O, 73.54% for Feature Set P, and 77.02% for Feature Set Q. This result is consistent with our results from the survival analyses in that early bids are informative of funding success.

In addition to funding success, we use early bids to predict the time it takes a listing to reach full funding (hours). As the funding duration (measured in hours) is a continuous outcome variable, we replace the Logistic model with Random Forest, while maintaining the use of LightGBM, XGBoost, and MLP, as these three algorithms can incorporate continuous outcome variable regressions. We also update our evaluation metrics to RMSE, MAE, MAPE, and R2, to better align with our shift from classification to regression analysis. From Table 7 we can see that, across all the algorithms, using Feature Set P (bids from the first funding day) significantly improves the prediction of funding time relative to Feature Set B (no RWOC bids).

[Insert Table 7 about here.]

Our results in this section demonstrate that adding early bids to predictive modeling can significantly improve the prediction performance of funding success and funding speed. The rationale behind this is intuitive. As previously illustrated, early, large, and above-normal bids are indicative of high-quality listings. If subsequent lenders interpret these bids in the same manner, such listings should be more likely to receive funding. We conduct several empirical tests to verify this mechanism and find that listings with more early large bids are likely to attract more subsequent lenders, larger bids, and be fully funded faster. This suggests that subsequent lenders are indeed capable of discerning the quality signals inherent in investment history and responding to them accurately. Details of these tests are in the Online Appendix IV.

7. RWOC Variables and Financial Inclusiveness

An original value proposition of the broader crowdfunding phenomenon is to better match the supply and demand of funds and to improve financial inclusiveness (G20 Global Partnership for Financial Inclusion white paper[15]). In this spirit, we examine the implications of RWOC variables for this important issue. Specifically, we investigate whether including RWOC variables would improve relative loan quality prediction and funding approval rate for borrowers who can more easily secure funding, relative to those who have more challenges in doing so. We focus on borrower differentiation by income, as low-income borrowers tend to have less financial access, and gender, as female borrowers tend to have more difficulty securing funds (Kanze et al. 2018).

For income, we classify all borrowers into high-income and low-income groups. We apply the trained algorithms using Feature Set B and Feature Set H from §5.1 to each category. We first find that the *increase* of AUC is larger for low-income borrowers than for their high-income counterparts (Table 8). Moreover, the increases of the true negative rate and the sum of the true negative rate and the false negative rate are both larger for low-income borrowers, suggesting a larger improvement in their funding approval rate and correctly predicted repaid rate. Hence, despite the higher risk associated with low-income borrowers, RWOC variables are useful for predicting their loan outcomes and therefore help investors find good candidates among them.

[Insert Table 8 about here.]

We similarly examine the benefits of RWOC variables for female versus male borrowers.[16] Table 8 shows that the *increase* of AUC, approval rate, and correctly predicted repayment rate of female

[15] https://www.gpfi.org/sites/gpfi/files/documents/GPFI_WhitePaper_Mar2016.pdf

[16] Prosper does not directly collect or reveal a borrower's gender. Hence, we encode the variable gender in two steps. We first analyze all the photos posted by borrowers through Microsoft Azure Face API, which is an AI service that analyzes faces in images. One of the analysis results returned from the API is the gender of the figure shown in an image. Second, we analyze borrower's narrative listing descriptions by keyword extraction. As Prosper suspended the display of personal photos in 2010 and the display of listing descriptions in 2013, we limit the listings in the sample of this test to 2005-2010. After these two steps, we identify 7,291 male borrowers and 9,444 female borrowers. We retrain Models A and B on the new sample and then apply the trained models on the male group and the female group separately.

borrowers are all larger than those of male borrowers. These results are consistent with recent findings that crowdfunding can help bridge gender gaps documented in traditional financing (Bapna and Ganco 2021).

Taken together, our results suggest that RWOC variables are particularly useful in "thin file" situations, where the financial information about borrowers is sparse. Traditional financial features tend to perform better at evaluating high-quality borrowers who have lower levels of uncertainty, which is also why such borrowers are typically categorized as high credit. These borrowers are more predictable in their future repayment behavior due to their established credit history and financial stability. In contrast, low-credit borrowers are generally harder to predict due to a lack of historical data or conflicting signals. The RWOC variables capture additional information about borrower quality that has not been captured by traditional financial features or credit scoring models, adding more marginal value to financially marginalized borrowers who have thinner credit files and lower baselines. This finding highlights the benefits of RWOC variables for further improving crowdfunding access for traditionally underserved groups of borrowers and aligns with the recent literature that suggests financial inclusion benefits from technological advancements (Li et al. 2024).

8. Discussion and Conclusions

Our paper studies the informational value of heterogeneous investment behavior in online debt crowdfunding. Using a combination of both descriptive and predictive models (Hofman et al. 2021, Shmueli and Koppius 2011), we propose simple metrics based on the amount (both absolute and relative) and timing of bids received from investors that improve the performance of loan quality predictive models. Under a descriptive framework based on an incremental classification of bids, we show that a simple count of the number of bids (Level 0) is not significantly associated with loan quality. By distinguishing minimum bids from the others, we find that only the number of bids that exceed the platform-mandated minimum is positively associated with loan quality. Further down the hierarchy, we first show that among bids that exceed the platform minimum, those that are large in absolute terms (defined as three standard deviations higher than the mean of all bids in the same week of the bid), those that are large in relative terms (compared to the focal investors' historical mode of bid amount), and large bids that are placed earlier in the funding period have increasingly stronger correlations with loan quality.

We then use a predictive modeling framework to demonstrate that incorporating RWOC variables improves loan quality predictions. We show that under a variety of predictive models, including those published in top journals based on data from the same context as well as different prediction algorithms, adding RWOC variables as additional predictors consistently improves predictive performance. Furthermore, we show that even if RWOC variables are derived only from bids placed in the first half of the funding period, they still predict loan quality well. They also predict the likelihood of full funding before they become loans. We also find that early bids contribute much more to predictive efficacy than later bids. Finally, we show the potential benefit of RWOC variables in promoting the inclusion of traditionally marginalized borrower groups.

The RWOC approach has several unique advantages. Regardless of what information investors can access, all that our approach requires is readily available to the platform administrators. RWOC variables represent detailed but still aggregated information about bid heterogeneity; moreover, calculations are easily implemented and scalable. It provides an efficient and effective manner of extracting crowd wisdom, superior to a simple aggregate of the crowd.

Our study makes several contributions to the theoretical and empirical research literature. First, even though the term "wisdom of the crowd" is used broadly in online communities and markets, the crowd is often viewed as a monolithic entity. Yet, no two individuals are truly identical, and actions by the same individual at different times may be driven by very different information. By recognizing that different actions have different values, we can better differentiate signals (informative actions) from noise (imitations or trivial actions). This mixture of useful and less-useful crowd actions is broadly applicable to other online communities and markets (e.g., crowdsourcing and open innovation), in which the crowd wisdom concept is applied.

Our study also contributes to the growing finance literature, especially in fintech, by proposing a new method of quality evaluation. While existing studies focus on borrower characteristics, our study highlights the possibilities of extracting information from investors' reactions. Whether the same logic applies to product markets remains an open question and is a fertile area for future research.

Our study also highlights a potential synergy between machine learning (predictive) and crowd wisdom. For example, rather than simply predicting loan quality based on borrower information, it allows learning from the reactions of the crowd. Future studies can and should examine how crowds react to such updated predictions.

Our study has multiple practical implications. First and foremost, our paper has important implications for platform design by providing an approach to distill several pieces of succinct information from the complicated list of bidding history in an easily scalable manner. Platforms can implement our method to help investors make decisions without potentially compromising privacy, especially when investors cannot directly access such data. P2P lending platforms such as Prosper.com removed bidding history to protect the privacy of investors. However, our analysis shows that there is considerable information in those bids that can help investors decide whether a loan is a good investment opportunity. By constructing and displaying RWOC-based variables (instead of the full bidding history), P2P lending platforms can continue to preserve the privacy of investors while also providing useful information from crowd reactions, so as to help investors make better decisions. For example, instead of showing the full bidding history, the platform can simply show one or more of the following metrics: (A) number of large bids based on typical bidding behaviors in the market; (B) number of bids in amounts that significantly exceed the typical bidding amount of their originating investors; and (C) number of bids in the first two categories that had been placed in the first day of the bidding process. These could be especially useful for retail investors who lack professional expertise.

For the primary market where the borrower requests P2P loans from lenders, our paper provides direct implications for the platform throughout the life cycle of the loan funding process. Before the fundraising process, if a borrower had requested (but failed) or received a loan on the platform (but not yet fully repaid), and now he or she wishes to request a loan again, the platform can further evaluate the risk of funding them (i.e., whether to put their request on the platform at all) by not just using their credit report information, but also how the crowd reacted in their previous requests using our RWOC-approach. Furthermore, during the fundraising process, RWOC variables can help investors make better decisions and invest in loans that are more likely to be repaid. This, in turn, helps protect the platform's reputation and protect investor interest. In fact, as our study shows, the RWOC variables can be especially helpful for loans that face more difficulty getting funded, such as those from female borrowers and those with lower credit grades. Also, RWOC variables based on early bids can help predict the likelihood that a loan would be fully funded. This is helpful for investors who wish to minimize their chances of investing in a loan that ends up not funded, as their funds will not generate interest before loans originate. Additionally, RWOC variables can help platforms create filtering mechanisms or recommendation systems to help investors decide which loans to invest in. Moreover, after funds are raised but before the loans are issued, our RWOC approach can also be easily applied as a checkpoint between full funding and loan issuance. This will help platforms avoid funding bad loans. Platforms always verify loans before funds are disbursed, even if a loan received 100% funding. At this stage, the platform could leverage the RWOC approach as an additional check for the loan's likelihood of default. It is possible (as our misclassification cost estimates show) that a loan could appear profitable if we only make predictions based on borrower and loan characteristics, but not so much when RWOC variables are considered.

Our approach can be immediately applied to evaluating P2P loan quality in secondary markets, in which individual or institutional investors sell shares of previously funded loans (Holden et al. 2020). Institutional investors increasingly participate in P2P lending platforms, such that P2P lending is now often known as "marketplace lending" instead of P2P lending. After the loans are funded, institutional investors who participate in P2P lending often need to package or sell those loans on secondary markets instead of holding them to maturity. The full bidding history information is readily available from the platform during that process, and RWOC variables can be easily constructed and help better price the loans for secondary market sales.

Our study also contributes to the ongoing discussion about tradeoffs between market efficiency and privacy. While more information typically improves efficiency, privacy is an increasingly significant concern. Further, each investor may or may not have the mental or computational capacity to process detailed bidding data. These reasons lead some platforms to stop providing detailed bidding histories. Our approach, especially analyzing only early bids, strikes a balance between these opposing needs by focusing on investor actions over investor identities. In lieu of providing detailed bidding histories, platforms could generate RWOC-based metrics behind a privacy wall and present them to investors.

In essence, our study proposes a novel mechanism to further mitigate the prevalent issue of information asymmetry in online platforms. The RWOC variables, if implemented, can help both platforms and investors improve their borrower screening capabilities. Platforms can create advanced filtering mechanisms or recommendation systems, guiding investors toward more viable loans, particularly in the early stages of funding, or use information from RWOC to determine whether to approve a loan at the end of the funding process. This helps ensure that creditworthy borrowers, especially those who are less appealing through traditional metrics, can have better chances of being funded. It will also help protect retail investors who may be more sensitive to risks, and improve overall market efficiency and effectiveness.

Our approach, at least conceptually speaking, can be applied beyond the context of debt-based crowdfunding. Future studies can evaluate the benefit of this approach (i.e., its generalizability) in other types of crowdfunding while taking contextual details into account. Equity crowdfunding will likely generate similar results if bids are not frequently withdrawn. Reward crowdfunding, however, may be less applicable due to the lack of flexibility in contribution amounts (i.e., a limited number of rewards). However, it may have implications for how fundraisers design their reward structure. Beyond crowdfunding, given the significance of secondary markets in our economic system (such as mortgage-backed securities), there will likely be value in assessing investor reactions to the underlying original loans in those markets as well.

Another interesting direction for future research is the threshold for participation in online financial markets. What is the smallest amount that investors should pay to participate in a loan, or a future in equity crowdfunding? Many markets establish low participation thresholds and lower them over time to increase the number of participants, which is an important metric for investors who evaluate these platforms as investment opportunities. Prosper.com initially required a $50 minimum but later reduced that to $25 before our study period. There are interesting variations among fintech startups. Among UK equity crowdfunding platforms, Crowdcube and Seedrs require as little as £10, whereas SyndicateRoom requires a minimum of £1000. These lower thresholds enable broader access but may also reduce bid informational value: minimum bids, in our study, were poor predictors of loan quality, regardless of how many there were. It remains an open question whether such changes would make a crowd more prone to irrationality.

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., ... & Zheng, X. (2016). TensorFlow: a system for large-scale machine learning. *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, 265-283.

Agrawal, A., Catalini, C., & Goldfarb, A. (2014). Some simple economics of crowdfunding. *Innovation Policy and the Economy*, *14*(1), 63-97.

Agrawal, A., Catalini, C., & Goldfarb, A. (2015). Crowdfunding: Geography, social networks, and the timing of investment decisions. *Journal of Economics & Management Strategy*, *24*(2), 253-274.

Atanasov, P., Rescober, P., Stone, E., Swift, S. A., Servan-Schreiber, E., Tetlock, P., & Mellers, B. (2017). Distilling the wisdom of crowds: Prediction markets vs. prediction polls. *Management Science*, *63*(3), 691-706.

Bansal, S., & Gutierrez, G. J. (2020). Estimating uncertainties using judgmental forecasts with expert heterogeneity. *Operations Research*, *68*(2), 363-380.

Bapna, R., Goes, P., & Gupta, A. (2003). Analysis and design of business-to-consumer online auctions. *Management Science*, *49*(1), 85-101.

Bapna, S., & Ganco, M. (2021). Gender gaps in equity crowdfunding: Evidence from a randomized field experiment. *Management Science*, *67*(5), 2679-2710.

Bayus, B. L. (2013). Crowdsourcing new product ideas over time: An analysis of the Dell IdeaStorm community. *Management Science*, *59*(1), 226-244.

Ben-Assuli, O., & Padman, R. (2020). Trajectories of repeated readmissions of chronic disease patients: Risk stratification, profiling, and prediction. *MIS Quarterly*, *44*(1), 201-227.

Benjamin, V., & Raghu, T. S. (2023). Augmenting social bot detection with crowd-generated labels. *Information Systems Research*, *34*(2), 487-507.

Berg, T., Burg, V., Gombović, A., & Puri, M. (2020). On the rise of Fintechs: Credit scoring using digital footprints. *The Review of Financial Studies*, *33*(7), 2845-2897.

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, 841-890.

Bikhchandani, S., Hirshleifer, D., Tamuz, O., & Welch, I. (2021). *Information cascades and social learning* (NBER Working Paper No. w28887). National Bureau of Economic Research. https://www.nber.org/papers/w28887.

Boudreau, K. J., Lacetera, N., & Lakhani, K. R. (2011). Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Science*, *57*(5), 843-863.

Bowers, A. (2015). Relative comparison and category membership: The case of equity analysts. *Organization Science*, *26*(2), 571-583.

Budescu, D. V., & Chen, E. (2015). Identifying expertise to extract the wisdom of crowds. *Management Science*, *61*(2), 267-280.

Burtch, G., Ghose, A., & Wattal, S. (2013). An empirical examination of the antecedents and consequences of contribution patterns in crowd-funded markets. *Information Systems Research*, *24*(3), 499-519.

Burtch, G., Ghose, A., & Wattal, S. (2014). Cultural differences and geography as determinants of online prosocial lending. *MIS Quarterly*, *38*(3), 773-794.

Burtch, G., Ghose, A., & Wattal, S. (2016). Secret admirers: An empirical examination of information hiding and contribution dynamics in online crowdfunding. *Information Systems Research*, *27*(3), 478-496.

Chen, H., De, P., Hu, Y. J., & Hwang, B. H. (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. *The Review of Financial Studies*, *27*(5), 1367-1403.

Chen, J., & Stallaert, J. (2014). An economic analysis of online advertising using behavioral targeting. *MIS Quarterly*, *38*(2), 429-450.

Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. New York, USA.

Choi, H., Mela, C. F., Balseiro, S. R., & Leary, A. (2020). Online display advertising markets: A literature review and future directions. *Information Systems Research, 31*(2), 556-575.

Cox, D. R. (1958). The regression analysis of binary sequences. *Journal of the Royal Statistical Society: Series B (Methodological), 20*(2), 215-232.

Da, Z., & Huang, X. (2020). Harnessing the wisdom of crowds. *Management Science*, 66(5), 1847-1867.

Dai, H., & Zhang, D. J. (2019). Prosocial goal pursuit in crowdfunding: Evidence from Kickstarter. *Journal of Marketing Research*, *56*(3), 498-517.

DeLong, E. R., DeLong, D. M., & Clarke-Pearson, D. L. (1988). Comparing the areas under two or more correlated receiver operating characteristic curves: A nonparametric approach. *Biometrics*, *44*(3): 837-845.

Donovan, J. (2021). Financial reporting and entrepreneurial finance: Evidence from equity crowdfunding. *Management Science*, *67*(11), 7214-7237.

Duarte, J., Siegel, S., & Young, L. (2012). Trust and credit: The role of appearance in peer-to-peer lending. *The Review of Financial Studies*, *25*(8), 2455-2484.

Ewens, M., & Townsend, R. R. (2020). Are early stage investors biased against women? *Journal of Financial Economics*, *135*(3), 653-677.

Frydman, C., & Krajbich, I. (2022). Using response times to infer others' private information: an application to information cascades. *Management Science*, *68*(4), 2970-2986.

Fu, R., Huang, Y., & Singh, P. V. (2021). Crowds, lending, machine, and bias. *Information Systems Research, 32*(1), 72-92.

Gao, Q., Lin, M., & Sias, R. W. (2023). Words matter: The role of texts in online credit markets. *Journal of Financial and Quantitative Analysis*, *58*(1), 1-28.

Greenstein, S., & Zhu, F. (2018). Do experts or crowd-based models produce more bias? Evidence from Encyclopedia Britannica and Wikipedia. *MIS Quarterly*, *42*(3), 945-959.

Grenadier, S. R., & Wang, N. (2005). Investment timing, agency, and information. *Journal of Financial Economics*, *75*(3), 493-533.

Hasan, I., He, Q., & Lu, H. (2022). Social capital, trusting, and trustworthiness: Evidence from peer-to-peer lending. *Journal of Financial and Quantitative Analysis*, *57*(4), 1409-1453.

Haykin, S. (1994). *Neural networks: a comprehensive foundation*. Prentice Hall PTR.

Herzenstein, M., Sonenshein, S., & Dholakia, U. M. (2011). Tell me a good story and I may lend you money: The role of narratives in peer-to-peer lending decisions. *Journal of Marketing Research, 48*(SPL), S138-S149.

Hildebrand, T., Puri, M., & Rocholl, J. (2017). Adverse incentives in crowdfunding. *Management Science*, *63*(3), 587-608.

Hofman, J. M., Watts, D. J., Athey, S., Garip, F., Griffiths, T. L., Kleinberg, J., ... & Yarkoni, T. (2021). Integrating explanation and prediction in computational social science. *Nature*, *595*(7866), 181-188.

Holden, C. W., Lin, M., Lu, K., Schneemeier, J., Wei, Z., & Yang, J. (2022). The Effect of Secondary Market Closure on Primary Market Funding Time: Evidence from Peer-to-Peer Lending. Working paper, Indiana University, Bloomington, Indiana.

Hong, Y., Hu, Y., & Burtch, G. (2018). Embeddedness, pro-sociality, and social influence: Evidence from online crowdfunding. *MIS Quarterly*, *42*(4), 1211-1224.

Iyer, R., Khwaja, A. I., Luttmer, E. F., & Shue, K. (2016). Screening peers softly: Inferring the quality of small borrowers. *Management Science*, *62*(6), 1554-1577.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An Introduction to Statistical Learning*. New York: Springer.

Jeppesen, L. B., & Lakhani, K. R. (2010). Marginality and problem-solving effectiveness in broadcast search. *Organization Science, 21*(5), 1016-1033.

Jiang, Y., Ho, Y. , Yan, X., & Tan, Y. (2020). When online lending meets real estate: Examining investment decisions in lending-based real estate crowdfunding. *Information Systems Research, 31*(3), 715-730.

Jiang, Y., Ho, Y. C., Yan, X., & Tan, Y. (2022). What's in a "username"? The effect of perceived anonymity on herding in crowdfunding. *Information Systems Research*, *33*(1), 1-17.

Kanze, D., Huang, L., Conley, M. A., & Higgins, E. T. (2018). We ask men to win and women not to lose: Closing the gender gap in startup funding. *Academy of Management Journal*, *61*(2), 586-614.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., … Liu, T.-Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems*, *30*, 3146–3154.

Kim, K., & Viswanathan, S. (2019). The Experts in the crowd: The role of experienced investors in a crowdfunding market. *MIS Quarterly*, *43*(2), 347-372.

Konovalov, A., & Krajbich, I. (2017). Revealed indifference: Using response times to infer preferences. Konovalov, A., & Krajbich, I. (2019) Revealed strength of preference: Inference from response times. *Judgment and Decision Making*, *14*(4), 381-394.

Krishnamurthi, L., & Raj, S. P. (1988). A model of brand choice and purchase quantity price sensitivities. *Marketing Science*, *7*(1), 1-20.

Lee, H. C. B., Ba, S., Li, X., & Stallaert, J. (2018). Salience bias in crowdsourcing contests. *Information Systems Research*, *29*(2), 401-418.

Li, C., Wang, H., Jiang, S., & Gu, B. (2024). The Effect of AI-Enabled Credit Scoring on Financial Inclusion: Evidence from 1M Underserved Population. *MIS Quarterly*, forthcoming.

Lin, M., & Viswanathan, S. (2016). Home bias in online investments: An empirical study of an online crowdfunding market. *Management Science*, *62*(5), 1393-1414.

Lin, M., Prabhala, N. R., & Viswanathan, S. (2013). Judging borrowers by the company they keep: Friendship networks and information asymmetry in online peer-to-peer lending. *Management Science*, *59*(1), 17-35.

Lin, M., Sias, R. W., & Wei, Z. (2023). Experts vs. non-experts in online markets. *MIS Quarterly*, 47(1), 97-126.

Liu, D., Brass, D. J., Lu, Y., & Chen, D. (2015). Friendships in online peer-to-peer lending: pipes, prisms, and relational herding. *MIS Quarterly*, *39*(3), 729-742.

Lukyanenko, R., Parsons, J., Wiersma, Y. F., & Maddah, M. (2019). Expecting the unexpected: Effects of data collection design choices on the quality of crowdsourced user-generated content. *MIS Quarterly*, *43*(2), 623-648.

Majchrzak, A., & Malhotra, A. (2016). Effect of knowledge-sharing trajectories on innovative outcomes in temporary online crowds. *Information Systems Research*, *27*(4), 685-703.

Mannes, A. E. (2009). Are we wise about the wisdom of crowds? The use of group judgments in belief revision. *Management Science*, *55*(8), 1267-1279.

Mollick, E., & Nanda, R. (2016). Wisdom or madness? Comparing crowds with expert evaluation in funding the arts. *Management Science*, *62*(6), 1533-1553.

Moritz, A., & Block, J. H. (2016). Crowdfunding: a literature review and research directions. *Crowdfunding in Europe*, 25-53.

Müller-Trede, J., Choshen-Hillel, S., Barneron, M., & Yaniv, I. (2018). The wisdom of crowds in matters of taste. *Management Science*, *64*(4), 1779-1803.

Netzer, O., Lemaire, A., & Herzenstein, M. (2019). When words sweat: Identifying signals for loan default in the text of loan applications. *Journal of Marketing Research*, *56*(6), 960-980.

Padmanabhan, B., Fang, X., Sahoo, N., & Burton-Jones, A. (2022). Machine learning in information systems research. *Management Information Systems Quarterly*, *46*(1), iii-xix.

Pang, M. S. (2017). Politics and information technology investments in the US Federal government in 2003-2016. *Information Systems Research*, *28*(1), 33-45.

Papke, L. E., & Wooldridge, J. M. (1996). Econometric methods for fractional response variables with an application to 401 (k) plan participation rates. *Journal of Applied Econometrics*, *11*(6), 619-632.
Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., … others. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.
Samuelson, P. A. (1938). A note on the pure theory of consumer's behaviour. *Economica*, *5*(17), 61-71.
Shmueli, G., & Koppius, O. R. (2011). Predictive analytics in information systems research. *MIS Quarterly*, *35*(3), 553-572.
Simester, D., Timoshenko, A., & Zoumpoulis, S. I. (2020). Targeting prospective customers: Robustness of machine-learning methods to typical data challenges. *Management Science*, *66*(6), 2495-2522.
Simpson, B., Schreier, M., Bitterl, S., & White, K. (2021). Making the world a better place: How crowdfunding increases consumer demand for social-good products. *Journal of Marketing Research*, *58*(2), 363-376.
Song, Y., Li, Z., & Sahoo, N. (2022). Matching returning donors to projects on philanthropic crowdfunding platforms. *Management Science*, *68*(1), 355-375.
Spann, M., & Skiera, B. (2003). Internet-based virtual stock markets for business forecasting. *Management Science*, *49*(10), 1310-1326.
Surowiecki, J. (2005). *The Wisdom of Crowds*. Anchor.
Trusov, M., Ma, L., & Jamal, Z. (2016). Crumbs of the cookie: User profiling in customer-base analysis and behavioral targeting. *Marketing Science*, *35*(3), 405-426.
Vallee, B., & Zeng, Y. (2019). Marketplace lending: A new banking paradigm? *The Review of Financial Studies*, *32*(5), 1939-1982.
Wang, G., Chen, G., Zhao, H., Zhang, F., Yang, S., & Lu, T. (2021). Leveraging Multisource Heterogeneous Data for Financial Risk Prediction: A Novel Hybrid-Strategy-Based Self-Adaptive Method. *MIS Quarterly*, *45*(4).
Wang, H., & Overby, E. (2020 January). How does algorithmic trading influence investor participation in peer-to-peer online lending markets? *Hawaii International Conference on System Sciences*, Maui, HI.
Wang, J., Ipeirotis, P. G., & Provost, F. (2017). Cost-effective quality assurance in crowd labeling. *Information Systems Research*, *28*(1), 137-158.
Wei, Z., & Lin, M. (2017). Market mechanisms in online peer-to-peer lending. *Management Science*, *63*(12), 4236-4257.
Younkin, P., & Kuppuswamy, V. (2018). The colorblind crowd? Founder race and performance in crowdfunding. *Management Science*, *64*(7), 3269-3287.
Zhang, J., & Liu, P. (2012). Rational herding in microloan markets. *Management Science*, *58*(5), 892-912.
Zheng, A., & Casari, A. (2018). *Feature engineering for machine learning: principles and techniques for data scientists*. O'Reilly Media, Inc..

Tables and Figures

Figure 1. Hierarchical Bid Classification

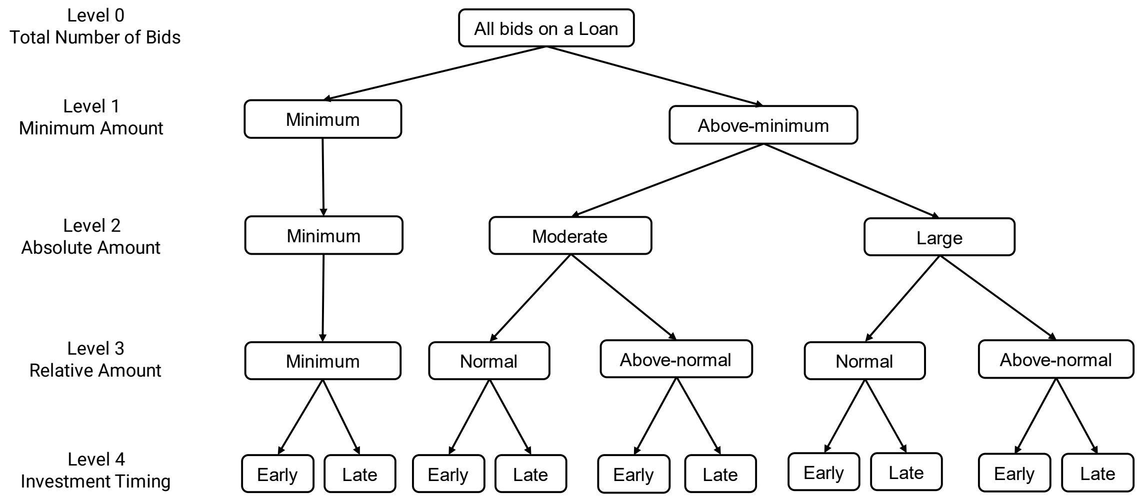

Notes. Level 1 separates bids into minimum and above minimum. Level 2 further classifies above-minimum bids as either moderate or large. Level 3 distinguishes moderate and large bids into normal or above-normal classes using individual lenders' histories. Level 4 divides each bid type in Level 3 into early and late classes. For each level, all bid types sum up to total bids.

Figure 2. Predictive Performance Increases with Funding Percentage and Funding Hours

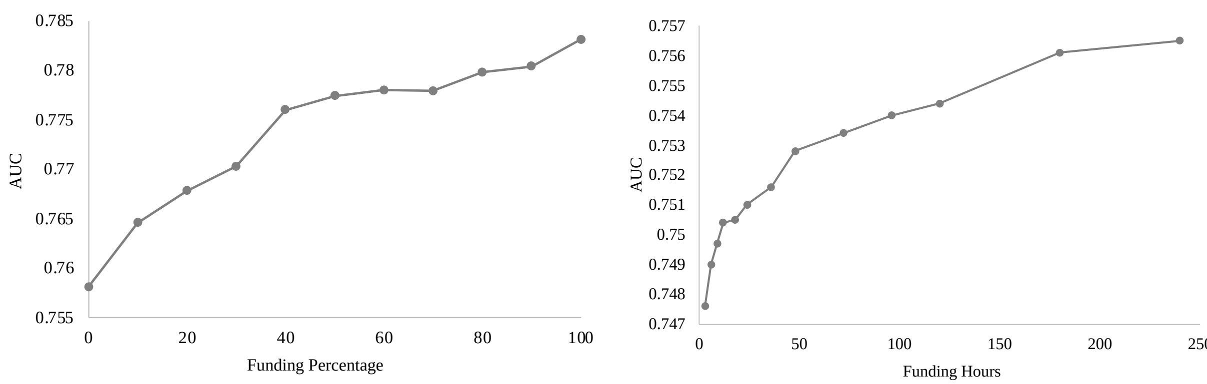

(A) Funding Percentage (B) Funding Hours

Notes. The predictive performance is based on the LightGBM algorithm. AUC increases as a function of funding percentage and funding hours at the beginning but slows down later.

Table 1. Descriptive Statistics

Variable name	Mean	SD	Min.	Max.	Median
Panel A: Loan quality					
Default percentage	0.14	0.31	0.00	1.00	0.00
Default indicator	0.23	0.42	0.00	1.00	0.00
ROI	1.15	0.39	0.00	2.13	1.21
Panel B: Variables about bid types					
#Bids	83.44	84.65	2.00	656.00	56.00
#Minimum bids	42.60	47.64	0.00	462.00	28.00
#Above-minimum bids	40.84	38.63	1.00	291.00	28.00
#Moderate bids	40.18	38.63	0.00	291.00	28.00
#Large bids	0.66	0.81	0.00	8.00	1.00
#Normal moderate bids	21.78	22.84	0.00	182.00	14.00
#Above-normal moderate bids	18.41	18.23	0.00	157.00	13.00
#Normal large bids	0.25	0.50	0.00	5.00	0.00
#Above-normal large bids	0.40	0.67	0.00	8.00	0.00
#Early minimum bids	15.13	16.46	0.00	319.00	11.00
#Late minimum bids	27.48	38.76	0.00	436.00	14.00
#Early normal moderate bids	8.55	8.74	0.00	118.00	6.00
#Late normal moderate bids	13.23	18.00	0.00	163.00	6.00
#Early above-normal moderate bids	7.19	7.84	0.00	121.00	5.00
#Late above-normal moderate bids	11.22	13.67	0.00	152.00	6.00
#Early normal large bids	0.08	0.28	0.00	4.00	0.00
#Late normal large bids	0.17	0.43	0.00	5.00	0.00
#Early above-normal large bids	0.13	0.38	0.00	5.00	0.00
#Late above-normal large bids	0.28	0.55	0.00	6.00	0.00
Panel C: Control variables					
Funded Amount (log)	8.70	0.67	7.60	10.13	8.59
Interest Rate	0.22	0.08	0.05	0.33	0.23
Estimated Return	0.11	0.03	0.02	0.18	0.12
Estimated Loss	0.09	0.05	0.00	0.20	0.09
Prosper Score	6.12	2.20	1.00	10.00	6.00
Now Delinquent	0.42	1.24	0.00	27.00	0.00
Amount Delinquent	1190.00	7726.00	0.00	279970.00	0.00
Public Records Last 12 Months	0.01	0.13	0.00	4.00	0.00
Public Records Last 10 Years	0.28	0.63	0.00	12.00	0.00
Delinquencies Last 7 Years	3.89	9.43	0.00	99.00	0.00
Inquiries Last 6 Months	1.13	1.58	0.00	27.00	1.00
Credit Age (Years)	17.28	7.78	0.62	61.22	16.25
Current Credit Lines	9.49	5.40	0.00	59.00	9.00
Open Credit Lines	8.42	4.87	0.00	48.00	8.00
Total Credit Lines	6.45	4.34	0.00	47.00	6.00
Revolving Credit Balance (USD in thousands)	19.45	36.58	0.00	1073.00	8.79
Bankcard Utilization	0.53	0.32	0.00	2.23	0.55
Debt/Income Ratio	0.70	28.55	0.00	4404.00	0.20
Employment Duration (Month)	96.44	93.11	0.00	755.00	67.00
Campaign Duration (Days)	4.81	4.27	1.00	15.00	3.00
Description Length (Words)	235.20	252.90	0.00	4703.00	167.00

Notes. The sample includes 29,900 loans funded between January 2011 and December 2012. The table omits the categorical control variables, including Category, Credit Grade, Consumer Credit Score, Day of Week, Employment, Group Member, Have Listing Description, Home Ownership, Income Range, Quarter, and Term. A detailed explanation of variables is presented in Online Appendix I.

Table 2. Bid Types and Loan Quality

Dependent variable: Default percentage	(1) Level 0 Total number of bids	(2) Level 1 Minimum bids	(3) Level 2 Absolute amount	(4) Level 3 Relative amount	(5) Level 4 Investment timing
#Bids	0.050 (0.165)				
#Minimum bids		2.096*** (0.534)	2.000*** (0.536)	0.942 (0.606)	
#Above-minimum bids		-2.613*** (0.696)			
#Moderate bids			-2.782*** (0.703)		
#Large bids			-28.431** (13.082)		
#Normal moderate bids				0.965 (1.245)	
#Above-normal moderate bids				-4.970*** (0.927)	
#Normal large bids				-23.804 (17.093)	
#Above-normal large bids				-36.983** (16.878)	
#Early minimum bids					4.081*** (1.188)
#Late minimum bids					-0.085 (0.742)
#Early normal moderate bids					0.871 (2.162)
#Late normal moderate bids					1.449 (1.556)
#Early above-normal moderate bids					-9.014*** (1.857)
#Late above-normal moderate bids					-3.714*** (1.287)
#Early normal large bids					-20.874 (28.543)
#Late normal large bids					-25.589 (19.916)
#Early above-normal large bids					-54.723** (25.489)
#Late above-normal large bids					-24.928 (19.077)
Observations	29,900	29,900	29,900	29,900	29,900
Loan information	YES	YES	YES	YES	YES
Borrower information	YES	YES	YES	YES	YES
Quarter dummies	YES	YES	YES	YES	YES

Notes. Coefficients are estimated using a fractional outcome regression model. For brevity, the table shows covariates only for variables of interest, with some covariates omitted. The number of bids is in the thousands. Robust standard errors are in parentheses. *p < 0.1; **p < 0.05; ***p < 0.01.

Table 3. Predictive Performance of Default Indicator

Algorithm	Feature Set	ROC-AUC	Misclassification Cost
Logistic	(A) Credit score	0.6272	$6,886,724
	(B) Baseline variables	0.7190	$6,648,399
	(C) Baseline + total number of bids	0.7209	$6,619,877
	(D) Baseline + simple RWOC variables	0.7384	$6,335,270
	(E) Baseline + engineered absolute bid amount variables	0.7628	$5,698,974
	(F) Baseline + engineered relative bid amount variables	0.7491	$6,291,413
	(G) Baseline + engineered bid timing variables	0.7363	$6,498,680
	(H) Baseline + all engineered RWOC variables	0.7654	$5,688,476
	(I) Baseline + no relative amount variables	0.7633	$5,749,112
	(J) All engineered RWOC variables	0.6955	$6,121,821
LightGBM	(A) Credit score	0.6222	$6,886,724
	(B) Baseline variables	0.7581	$5,456,783
	(C) Baseline + total number of bids	0.7552	$5,510,822
	(D) Baseline + simple RWOC variables	0.7643	$4,378,800
	(E) Baseline + engineered absolute bid amount variables	0.7824	$4,690,052
	(F) Baseline + engineered relative bid amount variables	0.7734	$4,778,969
	(G) Baseline + engineered bid timing variables	0.7635	$4,566,343
	(H) Baseline + all engineered RWOC variables	0.7831	$4,412,212
	(I) Baseline + no relative amount variables	0.7817	$4,248,388
	(J) All engineered RWOC variables	0.7318	$3,833,954
XGBoost	(A) Credit score	0.6251	$6,570,643
	(B) Baseline variables	0.7636	$5,159,066
	(C) Baseline + total number of bids	0.7624	$5,222,467
	(D) Baseline + simple RWOC variables	0.7672	$4,515,613
	(E) Baseline + engineered absolute bid amount variables	0.7832	$4,558,412
	(F) Baseline + engineered relative bid amount variables	0.7743	$4,440,922
	(G) Baseline + engineered bid timing variables	0.7687	$4,537,821
	(H) Baseline + all engineered RWOC variables	0.7840	$4,521,064
	(I) Baseline + no relative amount variables	0.7824	$4,754,126
	(J) All engineered RWOC variables	0.7314	$4,164,881
MLP	(A) Credit score	0.6296	$6,886,724
	(B) Baseline variables	0.7543	$5,998,029
	(C) Baseline + total number of bids	0.7497	$6,073,459
	(D) Baseline + simple RWOC variables	0.7662	$5,970,391
	(E) Baseline + engineered absolute bid amount variables	0.7887	$5,364,104
	(F) Baseline + engineered relative bid amount variables	0.7804	$5,928,951
	(G) Baseline + engineered bid timing variables	0.7649	$5,965,027
	(H) Baseline + all engineered RWOC variables	0.7932	$5,345,816
	(I) Baseline + no relative amount variables	0.7880	$5,414,947
	(J) All engineered RWOC variables	0.7188	$5,442,303

Notes. Each row displays the performance of an algorithm using the corresponding feature set as predictors. Performance is measured using the area under the Receiver Operating Characteristics Curve (ROC-AUC) and misclassification cost. The outcome variable is Default Indicator. Logistic, multinomial logistic regression; LightGBM, Light Gradient Boosting Machine; XGBoost, eXtreme Gradient Boosting; MLP, multilayer perceptron. Feature Set A includes Consumer Credit Score and Prosper Score. Feature Set B includes all variables in Panel C of Table 1 (finance variables + campaign features). Feature Set C includes all features in Feature Set B and the variable in Level 0 of Figure 1 (i.e., the total number of bids in a loan). Feature Set D includes all features in Feature Set B and all the variables in Level 4 of Figure 1. Feature Set E includes all features in Feature Set B and the engineered RWOC features related to the absolute bid amount. Feature Set F includes all features in Feature Set B and the engineered RWOC features related to the relative bid amount. Feature Set G includes all features in Feature Set B and the engineered RWOC features related to bid timing. Feature Set H includes all features in Feature Set B and all the engineered RWOC variables in Feature Sets E, F, and G. Feature Set I includes all features in Feature Set B and the engineered RWOC features in Feature Sets E and G. Feature Set J includes only the engineered RWOC variables. Please see §5.1.1 for more details about the feature sets.

Table 4. Applying Our Models to Datasets in Literature

		Fu et al. (2021)	Our algorithms with Feature Set B	Iyer et al. (2016)	Our algorithms with Feature Set B
Sample		Loans originated between Mar. 2007 – Oct. 2008		Loans posted between Feb. 12, 2007 – Oct. 16, 2008	
Model	Regression	/	0.8247	0.7103	0.8242
	LightGBM	/	0.8299	/	0.8303
	XGBoost	0.741	0.8335	/	0.8333
	MLP	/	0.8346	/	0.8317

Notes. This table compares the predictive performance (ROC-AUC) of our algorithms using Feature Set A and two published predictive models. The published models were trained on the same data source (Prosper) as ours but in different time windows. When trained on the same time windows, our baseline models produce improved results. The outcome variable is Default Indicator. Regression in Iyer et al. (2016), linear regression; Regression in our algorithms, multinomial logistic regression; LightGBM, Light Gradient Boosting Machine; XGBoost, eXtreme Gradient Boosting; MLP, multilayer perceptron.

Table 5. Predictive Performance of Models Using Bids from Experts vs. Non-Experts

Feature Set	Logistic	LightGBM	XGBoost	MLP
(K) Baseline + engineered absolute bid amount variables	0.7610	0.7745	0.7781	0.7931
(L) Baseline + engineered absolute bid amount variables from experts	0.7559	0.7720	0.7757	0.7889
(M) Baseline + engineered absolute bid amount variables from non-experts	0.7601	0.7775	0.7793	0.7900
(N) Baseline + engineered absolute bid amount variables from experts + engineered absolute bid amount variables from non-experts	0.7596	0.7745	0.7775	0.7878

Notes. Each row displays ROC-AUC produced by an algorithm using the corresponding feature set as predictors. The outcome variable is default indicator. Finance, finance variables about borrower credit and loan features; Logistic, multinomial logistic regression; LightGBM, Light Gradient Boosting Machine; XGBoost, eXtreme Gradient Boosting; MLP, multilayer perceptron. Feature sets are based on absolute bid amount (i.e., large, medium, minimum), because many engineered variables from expert bids are 0 if we move further down along the hierarchy in Figure 1.

Table 6. Predictions using Early Bids

Algorithm	Feature Set	Default Indicator	Funding Success
Logistic	(B) Baseline variables	0.7190	0.6958
	(O) Baseline + early bids	0.7589	0.7875
	(P) Baseline + first-day bids	0.7581	0.8046
	(Q) Baseline + first-two-day bids	0.7588	0.8183
LightGBM	(B) Baseline variables	0.7581	0.7326
	(O) Baseline + early bids	0.7747	0.8445
	(P) Baseline + first-day bids	0.7762	0.8323
	(Q) Baseline + first-two-day bids	0.7756	0.8452
XGBoost	(B) Baseline variables	0.7636	0.7354
	(O) Baseline + early bids	0.7779	0.8448
	(P) Baseline + first-day bids	0.7780	0.8342
	(Q) Baseline + first-two-day bids	0.7787	0.8457
MLP	(B) Baseline variables	0.7543	0.7270
	(O) Baseline + early bids	0.7873	0.8257
	(P) Baseline + first-day bids	0.7925	0.8398
	(Q) Baseline + first-two-day bids	0.7902	0.8466

Notes. Column 1 displays the ROC-AUC values for predicting Default Indicator, and Column 2 displays the ROC-AUC values for predicting Funding Indicator. Each row represents an algorithm using the corresponding feature set as predictors. Finance, finance variables about borrower credit and loan features; Logistic, multinomial logistic regression; LightGBM, Light Gradient Boosting Machine; XGBoost, eXtreme Gradient Boosting; MLP, multilayer perceptron. Feature Set O includes all variables in Feature Set B and engineered RWOC variables that only use early bids. Feature Set P includes all variables in Feature Set B and engineered RWOC variables that only use the bids placed on the first day. Feature Set Q includes all variables in Feature Set B and engineered RWOC variables that only use the bids placed on the first two days. See §6.1 for more details about the feature sets.

Table 7. Funding Speed Prediction Using Early Bids

Algorithm	Feature Set	RMSE	MAE	MAPE	R2
RF	(B) Baseline variables	96.4380	80.9726	0.9389	0.3053
	(P) Baseline + first-day bids	59.3959	43.4174	0.3451	0.7365
LightGBM	(B) Baseline variables	93.5027	77.1589	0.8735	0.3470
	(P) Baseline + first-day bids	54.6959	39.6133	0.3121	0.7765
XGBoost	(B) Baseline variables	95.5746	78.6876	0.8941	0.3177
	(P) Baseline + first-day bids	57.9373	42.5952	0.3478	0.7493
MLP	(B) Baseline variables	102.1942	87.4923	1.0067	0.2200
	(P) Baseline + first-day bids	61.3672	45.9165	0.3818	0.7187

Notes. Each row displays ROC-AUC from an algorithm using the corresponding feature set as predictors. The outcome variable is funding time in hours. Baseline variables, finance variables about borrower credit and loan features; RF: Random Forest; LightGBM, Light Gradient Boosting Machine; XGBoost, eXtreme Gradient Boosting; MLP, multilayer perceptron.

Table 8. Value of RWOC Variables for Financial Inclusiveness

		Predictive Performance (ROC-AUC)	Approval Rate (True Negative + False Negative)	Correctly Predicted Repaid Rate (True Negative)
Income	High Income	Increased 13.65pp	Increased 6.54pp	Increased 9.73pp
	Low Income	Increased 15.50pp	Increased 7.06pp	Increased 10.27pp
Gender	Male	Increased 3.27pp	Increased 1.36pp	Increased 2.06pp
	Female	Increased 3.52pp	Increased 2.37pp	Increased 2.50pp

Notes. This table reports the change of performance for loan default prediction within each subgroup from Feature Set B to Feature Set H. We report the results from LightGBM for brevity; results from other algorithms are qualitatively consistent. pp stands for percentage point.